IBM @server Cluster 1350 Release Notes (April 2005)

This document provides information about known limitations and restrictions that apply to the April 2005 IBM eServer Cluster 1350 release.

Known Limitations	Workaround
Due to SUSE LINUX Enterprise Server (SLES) Version 9 for x86_64-bit software compatibility problems with the on board video used in the IBM xSeries® 336, xSeries 346, and the HS20 (8843) blade server network installations using Cluster Systems Management (CSM) hang after the images are copied. This compatibility issue is resolved with SLES Version 9 SP1.	**Workaround:** To allow the installation to complete successfully, the SLES Version 9 x86_64 installation program (YaST) must be prevented from probing the system video. To prevent the system video from being probed, on the management node, complete the following steps: 1. After you have set up the SLES installation source, loop mount the installation root filesystem to the mount point of your choice, by entering the following command: **mount -o loop -r /csminstall/Linux/SLES/9/x86_64/GA/boot/root <mountpoint>** 2. To copy the files in the installation root filesystem into a temporary directory, enter the following commands: **mkdir <tmpdir>** **cd <mountpoint>** **find . -print | cpio -dump <tmpdir>** 3. To unmount the loop-mounted installation root filesystem, enter the following command: **umount <mountpoint>** 4. Edit the Display.ycp file in the installation root filesystem, by entering the following commands: **cd <tmpdir>/usr/share/YaST2/modules** **vim Display.ycp** 5. Comment out lines 373 through line 375 by adding two backslashes (//) to the beginning of each line, then; save the changes to the Display.ycp file. 6. Make sure that the yast2-core RPM is installed then compile the new version of Display.ycp (overwriting the old Display.ybc file) by entering the following command (the current directory should still be <tmpdir>/usr/share/YaST2/modules), **ycpc -c Display.ycp** 7. Make sure that the util-linux-2.12-72.20 RPM is installed then make a new installation root filesystem file. Enter the following commands: **cd /root** **mkfs -t cramfs <tmpdir> root.new** 8. To copy the new installation root filesystem file into the SLES CSM installation source directory, enter the following commands: **cp /root/root.new /csminstall/Linux/SLES/9/x86_64/GA/boot/root**

Known Limitations	Workaround
Because SLES Version 9 has a conflict with the BladeCenter™ media tray, one HS20 blade server (8843) per BladeCenter unit can hang during a network installation of SLES Version 9.	Prior to installation, for any BladeCenter unit that has the media tray owned by an HS20 blade server (8843), either change the owner to a blade server that is not an HS20 blade server (8843), or to a blade server that is performing the installation. If all blades servers are to be installed, the media tray owner can be set to **None** through the Remote Control feature of the management module Web interface.
PCI devices fail to function properly under certain circumstances on the eServer 326 node with certain Myrinet driver levels because there is a PCI resource management issue with the eServer 326 node.	To solve this issue, after loading the GM driver, enter the following command: **setpci -d 14c1: 56.b=48**
When an xSeries 346 node is running the Red Hat Enterprise Linux (RHEL) Release 3 Update 3 32-bit or 64-bit operating system, adapters installed into slot 3 can fail to operate. If these adapters are moved to slot 4, the problem abates. If ACPI is turned on, interrupt 19 can be assigned to the PCI slot 3. This causes a conflict which results in adapter drivers failing to load.	Open the **grub.conf** file and set the kernel command **acpi=noirq**, as a kernel-command line parameter.
When loading the drivers for the HS20 Fibre Channel expansion card the adapter fails to initialize because the world-wide port name for the adapter has to be in the zone when using an external Brocade switch before the driver is loaded.	Add the world-wide port name of each adapter port to the zone on the external Brocade Fibre Channel switch.
Due to an incompatibility between the HS20 blade server (8843) and the RHEL Release 3 Update 3 operating system during shutdown, the HS20 blade-server (8843) kernel panics on reboot when running the RHEL Release 3 Update 3 operating system. This problem has been fixed in the release of RHEL Release 3 Update 4 operating system.	To reboot the node after a kernel panic, use the CSM software interface to issue an rpower command (**rpower -n <nodename> reboot**).
The Myrinet Ethernet line card has Spanning-Tree protocol enabled for host ports which prevents nodes from reliably executing Dynamic Host Control Protocol (DHCP). This can cause the network boot to fail. The Myrinet Ethernet line card is not supported for connection to other Ethernet switches, and the Copper Pass-thru module does not support Serial-Over-LAN (SOL), and eth1 (the bottom NIC card) must be used for the boot NIC due to a SOL interruption when booting the network over eth1.	Do not use the Myrinet Ethernet line card for attachment of blade server Ethernet ports, or for the boot Ethernet port on non-blade nodes.
The PCI-X iSCSI adapter card is only supported for the RHEL Release 3 Update 3 32-bit operating system.	Do not use the PCI-X iSCSI adapter card with any operating system other than the RHEL Release 3 Update 3 32-bit operating system.

Known Limitations	Workaround
The Broadcom bcm5700 7.3.5 driver fails to compile on the RHEL Release 3 Update 3 operating system in the ppc64 environment. The makefile is configured for SLES Version 8 SP3 for the 2.4 kernel.	**Workaround:** You must make the following changes to the makefile: **--- Makefile 2004-06-22 19:25:24.000000000 -0400** **+++ Makefile.rhel3 2005-02-03 12:11:44.000000000 -0500** **@@ -84,7 +84,7 @@** **LD=/opt/cross/bin/powerpc64-linux-ld** **endif** **- CFLAGS+=-fno-strict-aliasing -fno-common -fomit-frame-pointer -msoft-float -pipe -mminimal-toc -fno-builtin+ CFLAGS+=-fno-strict-aliasing -fno-common -fomit-frame-pointer -msoft-float -pipe -mminimal-toc -fno-builtin -m64** **endif** **ifdef DBG** **@@ -101,7 +101,7 @@** **tar: bcm5700.tgz** **bcm5700.o: b57um.o b57proc.o tigon3.o autoneg.o 5701rls.o tcp_seg.o b57diag.o** **- $(LD) -i $^ -o $@** **+ $(LD) -melf64ppc -i $^ -o $@** **endif # ifeq ($(BCM_KVER),2.6)**
Due to an incompatibility with ServerRAID™ Manager Version 7.10.18 and certain JREs, a segmentation fault occurs when the ServeRAID agent is started on an x346 node running the RHEL Release 3 Update 3 32-bit operating system.	Use the IBM-Java2-1.4.2 JRE, which is included with CSM software. This allows the ServeRAID agent to run correctly in this environment.
Because the on board Ethernet, Adaptec SCSI controllers, and QLogic Fibre Channel adapters exhaust system PCI ROM space resources, the xSeries 346 node with two QLogic Fibre Channel adapters experiences poor performance, hang on boot, or become non-functional.	There are two resolutions for this known limitation (the second resolution is the preferred method): • Disable the boot ROM of the on board Broadcom Ethernet NICs after the node installation. This boot ROM can be re-enabled temporarily for network BIOS updates. • Install a ServeRAID 7k adapter. This takes up less PCI ROM space than the on board controller.
Recently the Myrinet GM builds have changed to reference include files under the /lib/modules directory structure. On the SLES Version 9 operating system this is not always properly configured. The make fails while compiling Myrinet GM driver version 2.1.5 on the SLES Version 9 operating system with default configuration settings.	Use the command **./configure --with-linux=/usr/src/linux** to allow the build to complete successfully (with a configured kernel source tree).
When the 2.4.21-ELsmp Athlon kernel version is installed, the Athlon SMP kernel configuration does not enable the 64 GB memory setting. The kernel sees only the first 4 GBs of memory after installing the RHEL Release 3 Update 3 operating system on an AMD Opteron eServer 326 server that has more than 4 GBs of memory.	Reinstall the operating system using the x86_64 port of the RHEL Release 3 Update 3 operating system.
The JS20 blade server does not successfully boot from a DHCP server which sends responses back as broadcast packets.	To avoid a problem with the JS20 blade server, install a DHCP server that sends unicast responses back, such as, the SLES Version 8 SP3 DHCP server.

Known Limitations	Workaround
Due to a problem with the BIOS build 15A, the serial output in Power On Self Test (POST) on the xSeries 336 server is very slow.	Make sure that the BIOS level is at 18A before turning on console redirection. To avoid waiting for the slow boot process to complete, remove the CMOS battery and touch the leads together for about 15 seconds to clear the CMOS settings. This will allow the system to boot at the normal speed.
You must install 32-bit X libraries or the FAStT MSJ install fails.	• Install XFree86-libs-4.3.0-68.EL.i386.rpm, by using the following command **--nodeps** • For PPC architecture, run the command, **/FAStTMSJ<ver>_install_ppc.bin -i silent -DSILENT_INSTALL_SET="QMSJ_LA_PPC"**. • For others, run the command, **./FAStTMSJ<ver>_install.bin -i silent -DSILENT_INSTALL_SET="QMSJ_LA"**.
Due to an incompatibility between both the xSeries 336 and xSeries 346 node BIOS and the pxelinux/memdisk files from Syslinux Version 3.07, exiting the PXE ROM causes the system to hang when booting an xSeries 336 or xSeries 346 node to the network and the pxelinux config file for it is configured to local boot, (either automatically or due to user input).	Down-level the versions of the pxelinux.0 and memdisk files on the tftpserver on the xSeries 336 and xSeries 346 boot nodes to those from an older version of Syslinux. The pxelinux.0 and memdisk files from Syslinux Version 2.10 does work with both the xSeries 336 and xSeries 346 nodes.
When running MPI applications on the SLES Version 9 32-bit operating system, NETDEV watchdog time-outs and temporary link-down events can occur on the first on board Ethernet interface on an xSeries 336 node and an xSeries 346 node. There is a known limitation when using TSO with the bcm5700.ko device driver v7.3.5 under the SLES Version 9 32-bit operating system.	Disable the TSO using the bcm5700.ko module load option *enable_tso=n,...* where <*n*> represents the value for each Ethernet interface using the bcm5700.ko module, and the value for the on board interfaces on both the xSeries 336 node and the xSeries 346 node must be zero (0).

First Edition (April 2005)

(1P) P/N: 59P4389